FogMQ: A Message Broker System for Enabling Distributed, Internet-Scale IoT Applications over Heterogeneous Cloud Platforms

Sherif Abdelwahab and Bechir Hamdaoui
Oregon State University, abdelwas,hamdaoui@eecs.orst.edu

arXiv:1610.00620v1 [cs.NI] 3 Oct 2016

***Abstract*—Excessive tail end-to-end latency occurs with conventional message brokers as a result of having massive numbers of geographically distributed devices communicate through a message broker. On the other hand, broker-less messaging systems, though ensure low latency, are highly dependent on the limitation of direct device-to-device (D2D) communication technologies, and cannot scale well as large numbers of resource-limited devices exchange messages. In this paper, we propose FogMQ, a cloud-based message broker system that overcomes the limitations of conventional systems by enabling autonomous discovery, self-deployment, and online migration of message brokers across heterogeneous cloud platforms. For each device, FogMQ provides a high capacity device cloning service that subscribes to device messages. The clones facilitate near-the-edge data analytics in resourceful cloud compute nodes. Clones in FogMQ apply Flock, an algorithm mimicking flocking-like behavior to allow clones to dynamically select and autonomously migrate to different heterogeneous cloud platforms in a distributed manner.**

***Index Terms*—Message brokering architecture, cloud computing resource allocation, distributed IoT applications.**

I. Introduction

Many large-scale applications are sensitive to latency as they rely on messaging sub-systems between geographically distributed devices and cloud services. Even if 10% of messages were delayed for longer than 150-300 ms, applications like remote-assisted surgery and real-time situation-awareness may not be feasible [1], [2]. A bounded tail end-to-end latency is a cornerstone for the realization of large-scale Internet of Things (IoT) applications near the network edge [3], [4].

When devices communicate through a middle message broker, successive packets queuing in multi-hop paths becomes a major source of latency. For example the average end-to-end latency of messages exchanged using a Redis broker [5] in a close amazon data-center is three times longer than deploying the same broker one-hop away from devices. Broker-less messaging using device-to-device communication does not necessarily solve the successive packets queuing problem. In IoT applications, a device communicates a large number of messages with many devices. Devices' limited processing and memory capacities become another major source of latency for large-scale distributed applications. Experiments show that direct device-to-device (D2D) messages can experience double the end-to-end latency when compared to brokering the messages through a one-hop away broker (see Section II).

When devices are cloned in a one-hop away cloudlet [6], a device's clone can provide message brokering service so that interacting devices can communicate with low latency while allowing them to offload their computation to resourceful nodes. Of course, communicating through a one-hop away clone may still cause long *tail end-to-end latency* - considering the 99-th percentile of computation plus communication latencies between clones/devices - when the broker service relays messages to distant devices. If a clone can measure: $1)$ messaging demand with other devices/clones, $2)$ the tail latency experienced by messages, and $3)$ the potential latency of other cloudlets/cloud platforms, clones can self-migrate between cloud platforms to always ensure a bounded weighted tail end-to-end latency. We show how autonomous clone migration can mimic birds flocking and prove that it is stable and achieves a tight latency that is $(1+\epsilon)-$far from optimal.

The use of cloudlets and dynamic service migration to solve latency problems are not new: Cloudlets [7] reduce the single-hop latency from 0.5-1 seconds to tens of milliseconds, and technologies like MobiScud and FollowMe [8], [9] migrate clones to sustain an average single-hop Round Trip Time (RTT) at nearly 10 ms. Such schemes struggle to make optimal migration decisions despite using central control units as they: adopt too constraining migration metric (average single-hop latency) and trigger migration only if devices locations change [10]. However, applications in fog computing [4] necessitate the deployment of inter-networking clones in heterogeneous platforms (cloudlet/clouds) without centralized administration. In this fog environment, clones communicate with several geo-distributed devices and other clones where the *tail weighted end-to-end latency* becomes the primary latency measure instead of the average RTT of a single-hop.

We propose FogMQ, a clone brokering system design that allows clones to self-discover and autonomously migrate to potential cloud hosting platforms according to self-measured weighted tail end-to-end latency, thereby stabilizing clone deployments and achieving low latency. FogMQ has four key features:

1. Reduces end-to-end delays that arise from multi-hop message queuing by deploying message brokers at the network edge, and while accounting for the devices' communication and relationship traffic patterns;
2. Ensures bounded message latency of IoT applications, outperforming conventional message brokers, like Redis and Nats;
3. Autonomously discovers and migrates to heterogeneous,

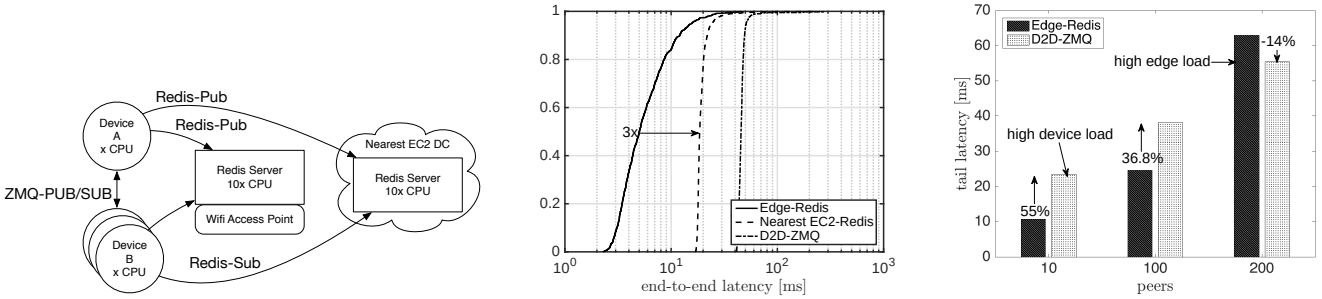

(a) Experiments setup. (b) CDF of latencies. (c) Tail latency in D2D vs edge brokers.

Fig. 1: Motivating experiments: the sources of latency in devices publish/subscribe communication.

Internet-scale cloud/edge platforms without needing centralized monitoring and control;

4. Simple and requires no change to existing cloud platforms controllers.

II. Motivation and Challenges

We first show that multi-hop queuing along Internet paths is a major source of end-to-end latency for IoT applications. We then show that devices' limited compute and memory resources standstill against latency reduction by direct device-to-device communication.

A. Sources of latency

When a message broker is hosted in a multi-tenant cloud platform, the end-to-end latency degrades due to *network interference*. Network interference occurs when the broker messages share: ingress/egress network I/O of its host, and one or more queues in the data-center network switches [11]. Hosts and network resources become spontaneously congested by traffic-demanding applications. For a single-authority cloud, an operator can control network interference of latency-sensitive applications with: switching, routing, and queuing management policies, besides controlling contention for hosts' compute, memory, and I/O resources [12], [13].

Network interference is harder to control for IoT applications. As devices communicate using our cloud-hosted broker, messages share network resources of multi-hop paths with diverse and unmonitored traffic. Multiple, unfederated authorities manage network resources along these paths, which makes it hard to enforce *unified* traffic shaping or queuing management polices. Adding also variations in devices' traffic demand, communication pattern with other devices, and mobility, it becomes particularity hard to trace devices' traffic, delays, and infrastructure conditions to find optimal policies with centralized solutions. Multi-hop queuing along Internet paths can account for a 3x degradation in end-to-end latency on average.

Fig. 1(a) illustrates an experiment setup to quantify this latency degradation. We install a Redis server in a Virtual Machine (VM)-instance in the nearest Amazon EC2 data-center (EC2-Redis), and install another Redis server in a same capacity VM in a host co-located with our WiFi access point (Edge-Redis). Our host runs other workloads. We emulate devices as simple processes running on another host that uses the same access point. We ensure that all VMs and hosts are time-synchronized with zero delay and jitter during the experiment execution time. A device emulator A publishes $10K$ messages to either Redis servers, and another device emulator B subscribes to A's messages. Fig. 1(b) shows the Cumulative Distribution Function (CDF) of the end-to-end latency, measured as the time between receiving a message at B and publishing it from A. The tail end-to-end latency for Edge-Redis is $15.6ms$, while it measured at $24.2ms$ for EC2-Redis accounting for 1.5x tail end-to-end latency improvement by avoiding multi-hop path to the closest EC2 instance and 3x improvement on average.

B. Why broker-less is not always the answer?

Benchmark	50%	$99^{th}\%$
Redis, 1000 messages	523.2	$1,276.0\mu s$
ZeroMQ, 1000 messages	314.8	$647.6\mu s$
Redis, $10,000$ messages	620.1	$2,010.5\mu s$
ZeroMQ, $10,000$ messages	320.3	$652.9\mu s$

TABLE I: Median and 99^{th} end-to-end latency of Redis and ZeroMQ measured under different loads (number of messages).

Direct device-to-device communication using broker-less message queues can thought of to be better than using message brokers. The obvious reasons for broker-less queues, such as ZeroMQ [14], superiority are: their lightweight implementation, and their usage of minimal number of shared queues, switches, routers, and access points, between communicating devices. TABLE I shows the median and tail end-to-end latency of ZeroMQ and Redis under different loads, where ZeroMQ can deliver $10,000$ messages three times faster than Redis.

Unfortunately, if the devices are resource limited, the latency superiority of direct device-to-device communication is not always maintained. Let us return to our motivating experiment in Fig. 1(a). We limit the resources used by the device emulators using Linux $\mathrm{cgroups}$ such that a device emulator can use no more than 10% of the CPU time compared to EC2-Redis or Edge-Redis. Fig. 1(b) shows that the average

end-to-end latency of D2D-ZeroMQ is 7 times longer than Edge-Redis, and the tail end-to-end latency is 4 times longer. Several factors can contribute to this deteriorated performance including the wireless environment loading and implementation details of either Redis or ZeroMQ. However, the main factor that limits direct device-to-device latency is the limited compute resources of the devices emulators.

To emphasis this observation, we increase the number of the publishing device peers (i.e. number of subscribing device emulators) until the Edge-Redis server becomes loaded. Fig. 1(c) shows the tail end-to-end latency for different numbers of peers. As we increase the number of peers, the latency superiority of the Edge-Redis starts to diminish, until we reach the 200 peers points at which our host becomes loaded at 90% utilization and the tail end-to-end latency of broker-less D2D-ZeroMQ becomes better by 14%. Broker-less device-to-device messaging is only better if a device computational resources are sufficiently large, which is an unrealistic assumption for most IoT devices.

III. FogMQ System

Our motivating experiments show that multi-hop queuing along Internet paths is a major source of tail end-to-end latency for cloud-based messaging systems, and that the latency improvement promise from device-to-device communication cannot always be attained due to limited devices resources. FogMQ tackles multi-hop queuing by reducing the queuing of messages: primarily, if message brokers can self-deploy and migrate in cloudlets according to the communication pattern of the devices, then the impact of multi-hop queuing delay can be diminished. In the extreme case, if two resource-limited devices communicate through brokers in the same cloudlet using the same access point, we can achieve a finite minimal bound on the latency. Autonomous brokers migration is the foundation idea of FogMQ.

In this section, we derive an intuitive design of FogMQ by which we bound weighted latency given an arbitrary network of heterogeneous cloud platforms. Although the stability and bounded performance of our design are intuitive, we solidify this intuition by relating the design to the theory of singleton weighted congestion games [15], [16], where we show that self-deploying clones reach a Nash Equilibrium (NE) and tightens the Price of Anarchy (PoA) of the weighted end-to-end delay.

A. Network of clones

To begin, we assume that devices communicate with each others according to IoT applications' requirements and form a social network of devices. Typically, the convergence of man-machine interactions in IoT will derive devices to form a social network [17]. This network can form according to existing social network structure of users or according to the required communication among devices that is inherited from application design.

The idea of application design for IoT is simple. An application is modeled as a graph (e.g. [18], [19], [20]). Each device participating in the execution of the IoT application *publishes* its data to its brokering clone. On the other hand, clones *subscribe* to each other according to the application-modeled graph, which forms an overlay network of clones to enable the IoT application. Upon completion of their executions, clones may *push* the results back to devices.

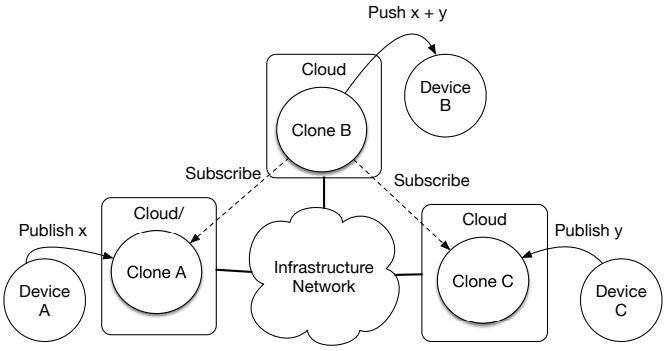

Fig. 2: Example overlay aggregation tree formed by device clones.

Fig. 2 illustrates a simple tree aggregation application for data retrieved from three devices. Device A and C publish their data x and y to their clones. As device B is interested in the result $x + y$, clone B subscribes to data from clones A and C to retrieve x and y, evaluates $x + y$, and pushes the result to its device. The advantages of using a pub/sub system for interacting between clones and the devices are: $1)$ providing an efficient messaging middleware to manage large-scale graph structures and multiple applications, $2)$ relying on the already in-place subscription and matching languages to effectively route information between devices and clones and inter-clones, and $3)$ simplifying the design of large-scale applications as overlay networks of and among the clones.

Generally speaking, the overlay network design of the clones is either structured or unstructured, and focuses mainly on minimizing a brokers fanout to minimize the communication between the clones. For example, topic-connected overlay networks are designed such that devices interested in the same topic are organized in a direct connected dissemination overlay [21]. The overlay network forms the foundation for distributed pub/sub, and directly impacts the system scalability and application performance [22], [23], [24]. We assume that an overlay topology of clones is given and we model it as a social network. We model the social network of clones as a graph $G = (V, P)$, where V denotes the set of n clones and P denotes the set of all clone pairs such that $p = (i, j) \in P$ if the i-th and j-th clones communicate with each other.

B. FogMQ architecture

We now describe how FogMQ initially creates device clones, as well as the FogMQ's architectural design tradeoffs. Fig. 3 illustrates the architectural elements of FogMQ.

FogMQ initially clones a device at the closest cloud/cloudlet from a set A of m cloud/cloudlets that are available to all devices and that can communicate over the Internet. An RPC client in the device is responsible for the clone initiation and

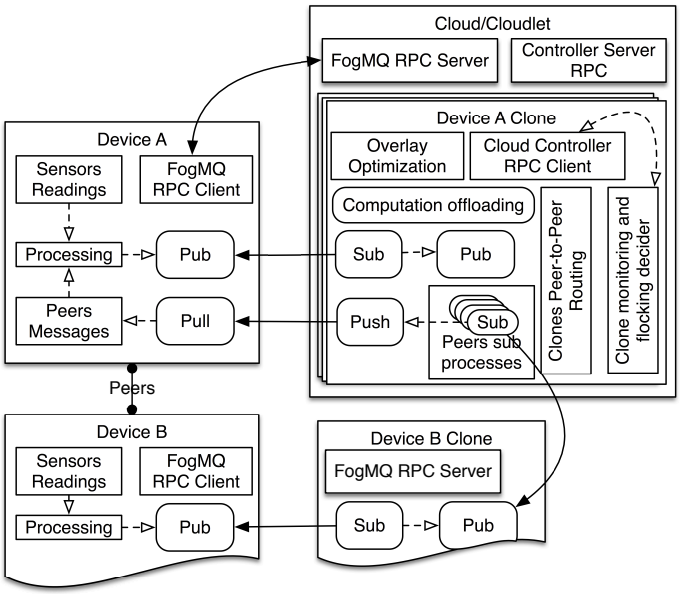

Fig. 3: FogMQ Architecture

peer relationship definition with other devices by which a device participates in the execution of a distributed application described as an overlay network of clones. Each cloud/cloudlet runs FogMQ RPC server as a middleware around the cloud/-cloudlet controller which implements different solutions that enable FogMQ to interact with heterogeneous cloud platforms (e.g. EC2, AWS, OpenStack cloudlet, or even a standalone host). A typical approach that clients can use to initiate clones is to query a global geo-aware domain name service load balancer to retrieve the IP address of the nearest FogMQ RPC server. With the integration of cloud computing in cellular systems [25], devices can also use native cellular procedures to initiate clones in the device's nearest cellular site.

The FogMQ RPC server realizes the device clone in a cloud platform as a virtual machine, container, or native process, where processes are always a favorable design choice to avoid latency overhead. Recent Linpackbenchmark [26] shows that containers and native processes can achieve a comparable number of floating point arithmetic per second, at least 2.5x greater than virtual machines. Despite that containers have better privacy and security advantage, as they provide a better administration, network, storage, and compute isolation, containers networking configuration can account for $30\mu s$ latency overhead when compared to that incurred by native processes [27]. As we will discuss later, implementing clones as processes has an advantage over both virtual machines and containers as they incur lesser migration overhead.

Once FogMQ creates a clone for a device, the clone subscribes to the devices' published messages that contain preprocessed sensors reading. Subscribing to the devices' messages eases clone migration processes as we will detail later. If a clone migrates from one cloud to the other, any changes to the clone IP address or network configuration become transparent to the device. Upon migration, the clone resubscribes to the devices' messages, allowing the device to continue publishing its messages without needing the clone to notify the device of its migration.

A clone can process its device's messages in its computation offloading module (see Fig. 3) with high processing, memory, and storage capacities. The computation offloading module of the clone also executes distributed applications defined as overlay networks that interconnect several clones. To exchange messages between peer-clones of an overlay network, a clone creates a separate process for each of its peers. Each process subscribes to published messages from its corresponding peer-clone to make messages from peers available for the computation offloading module. If needed a clone pushes messages and/or computation results back to its device using a push/pull messaging pattern. Fig. 3 illustrate messages flow between different modules for a simple example in which Device A is a peer to Device B.

The overlay optimization module and the peer-to-peer routing module are responsible for optimizing the fan-out of overlay networks and the routing decisions as we described earlier. Although the design of these mechanisms are integral to the performance of the overall system, the details overlay design and routing optimization algorithms are orthogonal to the scope of this paper in which we focus on autonomous migration decision that minimize the tail end-to-end latency.

C. Latency and peer-demand monitoring

Each device clone self-monitors and characterizes the demands with its peers, and evaluates latencies with the assistance of the hosting cloud. Let $d_{ij} \in \mathbb{R}^+$ denote the traffic demand between i and j and assume that $d_{ij} = d_{ji}$. Let $x_i \in A$ denote the cloud that hosts i and $l(x_i, x_j) > 0$ be the average latency between i and j when hosted at x_i and x_j, respectively (Note: if i and j are hosted at the same cloud, $x_i = x_j$). We assume that l is reciprocal and monotonic. Therefore, $l(x_i, x_j) = l(x_j, x_i)$ and there is an entirely nondecreasing order of $A \rightarrow A'$ such that for any consecutive $x_i, x_i' \in A'$, $l(x_i, x_j) \leq l(x_i', x_j)$. The reciprocity condition ensures that measured latencies are aligned with peer-VMs and imitates the alignment rule in bird flocking. We model $l(x_i, x_j) = \tau(x_i, x_j) + \rho(x_i) + \rho(x_j)$, where $\tau(x_i, x_j)$ is the average packet latency between x_i and x_j, and $\tau(x_i, x_j) = \tau(x_j, x_i)$. The quantity $\rho(x)$ is the average processing delay of x modeled as: $\rho(x) = \delta \sum_{i \in V: x_i = x} \sum_{j \in V} d_{ij} / (\gamma(x) - \sum_{i \in V: x_i = x} \sum_{j \in V} d_{ij})$, where δ is an arbitrary delay constant and $\gamma(x)$ denotes the capacity of x to handle all demanded traffic of its hosted VMs.

D. Learning new targets

Each cloud runs FogMQ RPC server as a middleware. The FogMQ servers in different clouds form a peer-to-peer network that evolves autonomously. Bootstrap nodes assist newly joined FogMQ servers to discover other servers. Gossip protocol is used to spread information about new FogMQ servers.

IV. Flock—An Adaptive Clone Migration Algorithm

Clones should adapt themselves to changes in the infrastructure network interconnecting the heterogenous cloud platforms, and should change according to the network state, structure, and applications' requirements. We propose an adaptive, fully distributed algorithm for dynamic cloud selection. The algorithm allows each VM i to learn a set $A_i \subseteq A$ (referred to as i's strategy set) from its hosting cloud x_i, and to autonomously select its hosting cloud based on local measurements only. Every cloud x updates its weight $w_x = \sum_{i:x_i=x} u_i(x)$ and broadcasts a monotonic, non-negative regularization function $f(w_x) : \mathbb{R}^+ \to \mathbb{R}^+$ with $\alpha < f(w_x) < 1$ for $\alpha > 0$ to each VM hosted at x. As VMs can each only be hosted by one cloud and all have access to the same set of strategies, we model the clone migration problem as a singleton symmetric weighted congestion game that minimizes the social cost $C(\sigma) = \sum_{x \in A} w_x f(w_x)$. If $f(w_x) \approx 1$, this game model approximates to minimizing $\sum_{i \in V} u_i(x_i)$. Let $x \xrightarrow{i} y$ denote that a VM i migrates from cloud x to cloud y and let $\eta \leq 1$ denote a design threshold. We now describe our proposed clone migration algorithm:

Flock: Autonomous VM migration protocol.

Initialization: Each clone $i \in V$ runs at a cloud $x \in A$.
Ensure: A Nash equilibrium outcome σ.

1: During round t, do in parallel: for all $i \in V$
2: i solicits its current set A_i from x.
3: i randomly selects a target cloud $y \in A_i$.
4: **if** $u_i(y) f(w_y + u_i(y)) \leq \eta u_i(x) f(w_x - u_i(x))$ **then**
5: $\quad x \xrightarrow{i} y$
6: **end if**

The following theoretical results have been proven in our work [28], and are included here for completeness.

Theorem 4.1: (Theorem 3.1 in [28]) Flock converges to a Nash equilibrium outcome.

Lemma 4.2: (Lemma 3.2 in [28]) The social value of Flock has a perfect PoA at most $\lambda/(1-\varepsilon)$ if for $\varepsilon < 1$ and $\lambda > 1-\varepsilon$ the regularization function satisfies $w^* f(w+w^*) \leq \lambda w^* f(w^*) + \varepsilon w f(w)$, where $w \geq 0$ and $w^* > 0$.

Theorem 4.3: (Theorem 3.3 in [28]) The regularization function $f(w) = exp(-1/(w+a))$ tightens the PoA to $1+\epsilon$ for a sufficiently large constant a and reduces the game to the original VM migration problem, i.e. minimizing $\sum_i u_i(x_i)$.

We now provide some simulation results, borrowed from [28] for completeness, to have an initial sense of how well Flock performs in terms of convergence and achievable PoA; more results on the performance of Flock can be found in [28]. Clouds are modelled as a complete graph with inter-cloud latency $\tau \sim$ Uniform$(10, 100)$ and cloud capacity $\gamma \sim$ Uniform$(50, 100)$. We model peer-to-peer clone relations as a binomial graph with $d \sim$ Uniform$(1, 10)$.

Fig. 4 shows the average number of rounds, k, needed for the algorithm to converge to a Nash equilibrium at 95%-confidence interval with 0.1 error. Observe that although the worst case of k is $O(n \log(n f_{max}))$ where f_{max} is the maximum value of the regularization function f [29], the figure shows that Flock scales better than $O(n)$ on average.

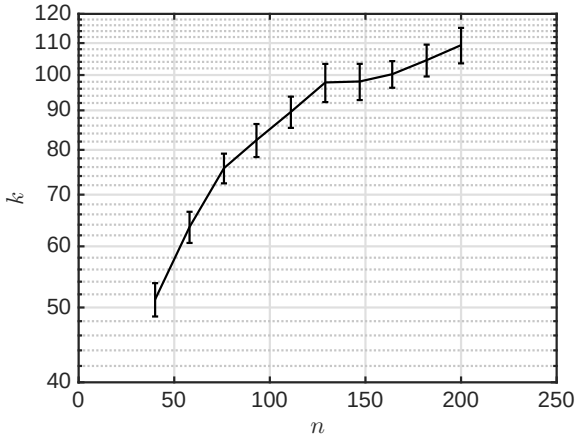

Fig. 4: Flock convergence: $a = 9$, $\eta = 0.9$, and # of clouds $= 37$.

Fig. 5 shows the PoA under different values of η. The optimal solution is implemented as brute-force in simulations to evaluate the minimum social value among all possible device clones-to-cloud assignments. Observe that as η approaches 1, the maximum achieved PoA matches the theoretical value of 1.21. In practical systems, η should not be too close to 1, so as to avoid migration oscillations, where clones keep migrating without much improvement.

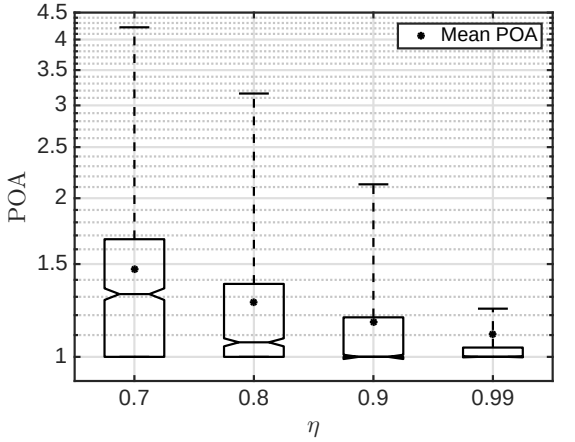

Fig. 5: $a = 9$, $n = 8$, and # of clouds $= 5$.

V. Related Work

We propose autonomous brokering clones for designing large-scale distributed pub/sub systems as a major mechanism that complements existing techniques to minimize the tail end-to-end messaging latency. Researches focus on three main techniques for the development of simple, scalable, and resource economic pub/sub systems: 1) content-centric above layer-3 routing between brokers (e.g. [30], [31], [32], [33], [34]), 2) overlay brokers network topologies designs (e.g.

[35], [36], [37], [23], [38]), and 3) content-centric in-network caching (e.g.[39], [40], [41], [41]).

Distributed pub/sub systems organize brokers, devices, or routing functions as an overlays and sub-overlays at the application layer. Upon constructing an efficient overlay network, routing protocols above layer-3 build minimum-cost message dissemination paths to deliver messages to subscribers according to specific topic-interest. Caching policies replicate clones' contents closer to devices interested in a content for faster repetitive publishing. For a given routing, overlay topology, and caching mechanisms, FogMQ ensures that these mechanisms achieve their full potentials by self-reorganizing the deployment of brokers through migrations in heterogeneous, unmanaged, and dynamic cloud environments. Unlike widely adopted centralized systems (e.g. Redis), FogMQ suits the large-scale applications and use cases of IoT and avoids the limitations of broker-less systems (e.g. ZeroMQ).

Existing migration solutions are limited in their applicability to minimize the weighted end-to-end latency. Several existing solutions rely on a system-wide central controller to manage the states of clones, devices, and physical resources of cloud platforms [10], [42]. For the considered fog environment, these solutions lack scalability for an Internet-sized network without relaxations that potentially compromise solutions quality.

Consider Markov Decision Process (MDP) based solutions. MDP requires a central server to collect statistics of devices mobility, clones demands, and clouds connectivity and utilization. This server also executes the value iteration algorithm to evaluate an optimal migration policy [10], [42], [43]. It is intractable to model all possible states of clones and their hosting platforms; hence it is common to discretize states measurements to relax the complexity of the policy optimization algorithms [43], [10]. This compromises the solutions quality.

Game-theoretic approaches potentially decentralize the migration algorithms and improve their scalability. However existing game-theoretic solutions provide an unbounded PoA [44]. We cannot use them - as they are - and guarantee optimal or close to optimal weighted end-to-end latency. Finally, existing migration solutions serve specialized cloud providers' objectives (e.g. energy, load, and cost) to profitably manage providers' infrastructures [45], [44]. *The existing models do not capture network latency between clones that are executing distributed IoT applications.* Unlike existing solutions, FogMQ adopts a simple autonomous migration protocol that is stable and bounds the tail end-to-end latency $(1-\epsilon)$ far from optimal.

VI. Conclusion and Future Work

We proposed FogMQ, a cloud-based, message broker system, composed of an architecture and an online migration algorithm, that enables autonomous discovery, self-deployment, and online migration of message brokers across heterogeneous cloud platforms. The migration algorithm, called Flock, enables autonomous discovery of and migration to heterogeneous cloud/edge platforms in a (i) decentralized manner and (ii) without requiring changes to existing cloud platform controllers. The proposed architecture enables the deployment of message brokers (i) at the edge clouds (i.e., cloudlets) near the end-user devices, and (ii) while accounting for the devices' communication and relationship traffic patterns.

An implementation of FogMQ on real cloud platforms is currently underway. In this implementation, clones are implemented as processes, Redis key-value store as a device registry, and edge clouds as Linux VMs. Our implementation-based performance evaluation of our proposed FogMQ system, a clone-based architecture design and an online clone migration algorithm, will be published when available.

Acknowledgement

This work was made possible by NPRP grant # NPRP 5-319-2-121 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

References

[1] M. Anvari, T. Broderick, H. Stein, T. Chapman, M. Ghodoussi, D. W. Birch, C. Mckinley, P. Trudeau, S. Dutta, and C. H. Goldsmith, "The impact of latency on surgical precision and task completion during robotic-assisted remote telepresence surgery," *Computer Aided Surgery*, vol. 10, no. 2, pp. 93–99, 2005.

[2] U. Ramachandran, K. Hong, L. Iftode, R. Jain, R. Kumar, K. Rothermel, J. Shin, and R. Sivakumar, "Large-scale situation awareness with camera networks and multimodal sensing," *Proceedings of the IEEE*, vol. 100, no. 4, pp. 878–892, 2012.

[3] S. M. Rumble, D. Ongaro, R. Stutsman, M. Rosenblum, and J. K. Ousterhout, "It's time for low latency." in *HotOS*, vol. 13, 2011, pp. 11–11.

[4] F. Bonomi, R. Milito, P. Natarajan, and J. Zhu, "Fog computing: A platform for internet of things and analytics," in *Big Data and Internet of Things: A Roadmap for Smart Environments*. Springer, 2014, pp. 169–186.

[5] J. L. Carlson, *Redis in Action*. Manning Publications Co., 2013.

[6] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *Pervasive Computing, IEEE*, vol. 8, no. 4, pp. 14–23, 2009.

[7] M. Satyanarayanan, Z. Chen, K. Ha, W. Hu, W. Richter, and P. Pillai, "Cloudlets: at the leading edge of mobile-cloud convergence," in *Mobile Computing, Applications and Services (MobiCASE), 2014 6th International Conference on*. IEEE, 2014, pp. 1–9.

[8] K. Wang, M. Shen, J. Cho, A. Banerjee, J. Van der Merwe, and K. Webb, "Mobiscud: A fast moving personal cloud in the mobile network," in *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*. ACM, 2015, pp. 19–24.

[9] T. Taleb and A. Ksentini, "Follow me cloud: interworking federated clouds and distributed mobile networks," *Network, IEEE*, vol. 27, no. 5, pp. 12–19, 2013.

[10] R. Urgaonkar, S. Wang, T. He, M. Zafer, K. Chan, and K. K. Leung, "Dynamic service migration and workload scheduling in edge-clouds," *Performance Evaluation*, vol. 91, pp. 205–228, 2015.

[11] S. K. Barker and P. Shenoy, "Empirical evaluation of latency-sensitive application performance in the cloud," in *Proceedings of the first annual ACM SIGMM conference on Multimedia systems*. ACM, 2010, pp. 35–46.

[12] X. Pu, L. Liu, Y. Mei, S. Sivathanu, Y. Koh, C. Pu, and Y. Cao, "Who is your neighbor: Net i/o performance interference in virtualized clouds," *Services Computing, IEEE Transactions on*, vol. 6, no. 3, pp. 314–329, 2013.

[13] R. C. Chiang and H. H. Huang, "Tracon: interference-aware scheduling for data-intensive applications in virtualized environments," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 47.

[14] P. Hintjens, *ZeroMQ: Messaging for Many Applications*. " O'Reilly Media, Inc.", 2013.

[15] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis, "The structure and complexity of nash equilibria for a selfish routing game," in *Automata, Languages and Programming*. Springer, 2002, pp. 123–134.

[16] D. Fotakis, S. Kontogiannis, and P. Spirakis, "Selfish unsplittable flows," *Theoretical Computer Science*, vol. 348, no. 2, pp. 226–239, 2005.

[17] L. Atzori, A. Iera, and G. Morabito, "Siot: Giving a social structure to the internet of things," *Communications Letters, IEEE*, vol. 15, no. 11, pp. 1193–1195, 2011.

[18] K. Hong, D. Lillethun, U. Ramachandran, B. Ottenwälder, and B. Koldehofe, "Mobile fog: A programming model for large-scale applications on the internet of things," in *Proceedings of the second ACM SIGCOMM workshop on Mobile cloud computing*. ACM, 2013, pp. 15–20.

[19] N. K. Ahmed, N. Duffield, J. Neville, and R. Kompella, "Graph sample and hold: A framework for big-graph analytics," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2014, pp. 1446–1455.

[20] N. Satish, N. Sundaram, M. M. A. Patwary, J. Seo, J. Park, M. A. Hassaan, S. Sengupta, Z. Yin, and P. Dubey, "Navigating the maze of graph analytics frameworks using massive graph datasets," in *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 2014, pp. 979–990.

[21] C. Chen, R. Vitenberg, and H.-A. Jacobsen, "A generalized algorithm for publish/subscribe overlay design and its fast implementation," in *Distributed Computing*. Springer, 2012, pp. 76–90.

[22] R. Baldoni, R. Beraldi, L. Querzoni, and A. Virgillito, "Efficient publish/subscribe through a self-organizing broker overlay and its application to siena," *The Computer Journal*, vol. 50, no. 4, pp. 444–459, 2007.

[23] C. Chen, H.-A. Jacobsen, and R. Vitenberg, "Algorithms based on divide and conquer for topic-based publish/subscribe overlay design," *IEEE/ACM Transactions on Networking*, 2015.

[24] Y. Sun, X. Qiao, B. Cheng, and J. Chen, "A low-delay, lightweight publish/subscribe architecture for delay-sensitive iot services," in *Web Services (ICWS), 2013 IEEE 20th International Conference on*. IEEE, 2013, pp. 179–186.

[25] S. Abdelwahab, B. Hamdaoui, M. Guizani, and T. Znati, "Replisom: Disciplined tiny memory replication for massive iot devices in lte edge cloud," 2015.

[26] J. J. Dongarra, P. Luszczek, and A. Petitet, "The linpack benchmark: past, present and future," *Concurrency and Computation: practice and experience*, vol. 15, no. 9, pp. 803–820, 2003.

[27] W. Felter, A. Ferreira, R. Rajamony, and J. Rubio, "An updated performance comparison of virtual machines and linux containers," in *Performance Analysis of Systems and Software (ISPASS), 2015 IEEE International Symposium On*. IEEE, 2015, pp. 171–172.

[28] S. Abdelwahab and B. Hamdaoui, "Flocking virtual machines in quest for responsive IoT cloud services," *arXiv preprint*, 2016.

[29] S. Chien and A. Sinclair, "Convergence to approximate nash equilibria in congestion games," in *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2007, pp. 169–178.

[30] C. Cañas, E. Pacheco, B. Kemme, J. Kienzle, and H.-A. Jacobsen, "Graps: A graph publish/subscribe middleware," in *Proceedings of the 16th Annual Middleware Conference*. ACM, 2015, pp. 1–12.

[31] N. Carvalho, F. Araujo, and L. Rodrigues, "Scalable qos-based event routing in publish-subscribe systems," in *Network Computing and Applications, Fourth IEEE International Symposium on*. IEEE, 2005, pp. 101–108.

[32] K. An, A. Gokhale, S. Tambe, and T. Kuroda, "Wide area network-scale discovery and data dissemination in data-centric publish/subscribe systems," *Network*, vol. 1, no. P4P3, p. P3, 2015.

[33] C. Chen and Y. Tock, "Design of routing protocols and overlay topologies for topic-based publish/subscribe on small-world networks," in *Proceedings of the Industrial Track of the 16th International Middleware Conference*. ACM, 2015, p. 2.

[34] W. Rao, K. Zhao, Y. Zhang, P. Hui, and S. Tarkoma, "Towards maximizing timely content delivery in delay tolerant networks," *Mobile Computing, IEEE Transactions on*, vol. 14, no. 4, pp. 755–769, 2015.

[35] G. Siegemund, V. Turau, and K. Maamra, "A self-stabilizing publish/subscribe middleware for wireless sensor networks," in *Networked Systems (NetSys), 2015 International Conference and Workshops on*. IEEE, 2015, pp. 1–8.

[36] C. Chen, Y. Tock, H.-A. Jacobsen, and R. Vitenberg, "Weighted overlay design for topic-based publish/subscribe systems on geo-distributed data centers," in *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*. IEEE, 2015, pp. 474–485.

[37] Y. Teranishi, R. Banno, and T. Akiyama, "Scalable and locality-aware distributed topic-based pub/sub messaging for iot," in *2015 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2015, pp. 1–7.

[38] M. Onus and A. W. Richa, "Minimum maximum-degree publish-subscribe overlay network design," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 5, pp. 1331–1343, 2011.

[39] M. Diallo, S. Fdida, V. Sourlas, P. Flegkas, and L. Tassiulas, "Leveraging caching for internet-scale content-based publish/subscribe networks," in *Communications (ICC), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1–5.

[40] K. Cho, M. Lee, K. Park, T. T. Kwon, Y. Choi, and S. Pack, "Wave: Popularity-based and collaborative in-network caching for content-oriented networks," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*. IEEE, 2012, pp. 316–321.

[41] M. Matos, A. Nunes, R. Oliveira, and J. Pereira, "Stan: exploiting shared interests without disclosing them in gossip-based publish/subscribe." in *IPTPS*, 2010, p. 9.

[42] V. Eramo, E. Miucci, and M. Ammar, "Study of migration policies in energy-aware virtual router networks," *Communications Letters, IEEE*, vol. 18, no. 11, pp. 1919–1922, 2014.

[43] L. Chen, H. Shen, and K. Sapra, "Distributed autonomous virtual resource management in datacenters using finite-markov decision process," in *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 2014, pp. 1–13.

[44] Z. Xiao, J. Jiang, Y. Zhu, Z. Ming, S. Zhong, and S. Cai, "A solution of dynamic vms placement problem for energy consumption optimization based on evolutionary game theory," *Journal of Systems and Software*, vol. 101, pp. 260–272, 2015.

[45] T. Duong-Ba, T. Nguyen, B. Bose, and T. Tran, "Joint virtual machine placement and migration scheme for datacenters," in *Global Communications Conference (GLOBECOM), 2014 IEEE*. IEEE, 2014, pp. 2320–2325.